Mon. Not. R. Astron. Soc. **000**, 1–12 (2016) Printed 2 December 2016 (MN LATEX style file v2.2)

A Joint Analysis for Cosmology and Photometric Redshift Calibration Using Cross Correlations

Michael McLeod[1⋆], Sreekumar T. Balan[1], Filipe B. Abdalla[1,2]
[1] *Department of Physics and Astronomy, University College London, Gower Place, London WC1E 6BT, U.K.*
[2] *Department of Physics and Electronics, Rhodes University, PO Box 94, Grahamstown, 6140, South Africa*

arXiv:1612.00307v1 [astro-ph.CO] 1 Dec 2016

2 December 2016

ABSTRACT

We present a method of calibrating the properties of photometric redshift bins as part of a larger Markov Chain Monte Carlo (MCMC) analysis for the inference of cosmological parameters. The redshift bins are characterised by their mean and variance, which are varied as free parameters and marginalised over when obtaining the cosmological parameters. We demonstrate that the likelihood function for cross-correlations in an angular power spectrum framework tightly constrains the properties of bins such that they may be well determined, reducing their influence on cosmological parameters and avoiding the bias from poorly estimated redshift distributions. We demonstrate that even with only three photometric and three spectroscopic bins, we can recover accurate estimates of the mean redshift of a bin to within $\Delta\mu \approx 3 - 4 \times 10^{-3}$ and the width of the bin to $\Delta\sigma \approx 1 \times 10^{-3}$ for galaxies near $z = 1$. This indicates that we may be able to bring down the photometric redshift errors to a level which is in line with the requirements for the next generation of cosmological experiments.

Key words: cosmology: large-scale structure - surveys

1 INTRODUCTION

Galaxy surveys have become in recent years an important source of data for cosmology, particularly for late time effects such as dark energy. Calculations of predicted statistical properties for a given cosmological model require the redshift distribution of observed galaxies to be known accurately. Spectroscopy has long been used to calculate accurate redshifts for objects, but this is a time intensive process requiring both detailed observation across the object's spectrum and careful analysis. In order to collect data for the vast numbers of galaxies required, current and future surveys are necessarily dependent on photometric redshifts for the majority of objects.

When photometry is used, approximate redshifts are calculated from a small number of intensities measured in (typically around five) broad bands. Standard methods of inferring redshifts from photometric data are to use machine learning methods, such as artificial neural networks, or to fit template functions (see e.g. Zheng & Zhang 2012 for a review). These require that we have large training sets of galaxies for which we have spectroscopic redshifts. Additionally, the spectroscopic set must be representative of the full photometric set (in terms of both redshift range *and* the nature of the objects contained within the sample) in order to reduce both the error and the bias in the derived relation (Sadeh et al. 2015, Abdalla et al. 2008). Unfortunately, the spectroscopic sample is rarely as large as we would like, and is even less often fully representative of the redshift range we wish to look at. Spectroscopic samples tend to be dominated by bright objects which are easier to study; a lack of spectroscopic objects, particularly at the extremes of the redshift range, tends to lead to larger errors in the redshift distributions reconstructed from machine learning techniques. (Often the middle of the range is reconstructed comparatively well, while the outer regions suffer.) Other methods may be more successful at the lower or higher end of redshift, but few methods can be confidently used across the entire range (Abdalla et al. 2008). In addition to this, the errors associated with any such reconstructions are large, and those who have attempted to reconstruct redshifts from photometric data will be familiar with the significant scatter around the spectroscopic redshifts (e.g. Banerji et al. 2008).

Due to the significant inherent uncertainties in such redshift estimates, for analysis objects may be grouped into bins of similar redshift. In order to achieve the precisions desired for current and future generations of cosmological experiments, we need to be able to determine the redshift distribution of each of these bins with greater accuracy than has been possible by simply using standard fitting to spectroscopic data. The impact of the redshift distributions on cosmology, and the importance of knowing them to a high

⋆ Email: michael.mcleod.13@ucl.ac.uk

degree of accuracy, has been the subject of a number of studies such as Huterer et al. (2004), Newman (2008). These suggest that to achieve the desired precision and accuracy in upcoming experiments such as LSST, we require that the mean and width of redshift bins to be known to $O(10^{-3}(1+z))$. Newman (2008) proposes that this may be achieved by calibrating the the photometric redshifts using information from cross-correlations with spectroscopic data.

Since the proposal of these ideas, there have been some studies looking into the potential for using cross-correlations for estimating photometric redshift distributions (e.g. Menard et al. (2014), Schmidt (2006), Schulz (2015), Matthews & Newman (1982), McQuinn & White (2013)), as well as potential problems such as contamination as in Benjamin et al. (2010). These tend to focus on recovering the redshift distribution from simulations by comparing the correlation between some photometric data set and a spectroscopic sample at known redshift, and assuming some fixed cosmology. In the case of a practical analysis however, we will not know the cosmological parameters (the determination of which is, after all, the objective of such calculations), and the calculation of theoretical correlation functions is cosmology dependent. It is well known that the cosmological parameters and redshift distribution are degenerate, and hence we cannot estimate how well the redshift distribution can be constrained without also varying the cosmology itself, as uncertainties from the cosmology may become a significant factor. This may be particularly important where the region of overlap between the photometric sample and the spectroscopic sample is relatively small. Hence, in order to avoid biases or overly optimistic estimates of our constraining power, we must determine the cosmology and the redshift distribution together, rather than treating them as independent problems. Previous work such as Newman (2008) and McQuinn & White (2013) also use estimators which may be prone to finding local maxima, and do not explore the space as fully as an MCMC approach using a full likelihood.

A significant amount of attention in recent years has been placed on the power of cross-correlations as a statistical tool for cosmology (such as Rhodes et al. (2014), Kirk et al. (2013)). Using cross-correlations to calculate both redshift distributions and cosmological parameters implies that we may include these effects into one framework with relative ease. We demonstrate such a technique for calibrating the photometric redshift distribution from an initial estimate using a joint likelihood analysis with cosmology using the angular power spectrum $C(l)$. We take the errors in photometric redshift modelling into account by allowing the mean and width of the photo-z bins to vary as free parameters, just as we do with cosmological parameters. The redshift binning is then marginalised over in order to obtain the probability contours for cosmological parameters. This allows us to study the extent to which we can constrain the photometric redshift distributions and simultaneously explore the impact of this information on cosmological inferences, in a manner which automatically treats the errors in the distributions in a Bayesian way.

In this paper we investigate extent to which photometric redshift bins can be constrained by cross correlations, and the impact on cosmological parameter inference in the case of large scale structure. We present a simplified experiment where we vary the width and mean of gaussian redshift bins, although we explain how the framework may be applied to higher moments also. For computational simplicity, only three photometric and three spectroscopic bins will be used (although this can be extended to fuller surveys at the expense of computation time); this simple model should suffice to demonstrate the power of the technique, as well as the degeneracies between the parameterisations of the redshifts and the cosmological parameters. The impact on future optical surveys will be greater though, as the same technique can be used to constrain photometric samples in weak lensing analyses, which may be used in conjunction with galaxy number counts to infer cosmology.

2 THE C(L) CALCULATION

The angular power spectrum is split up into correlations between different bins and cosmological probes; the full object we wish to look at is $C^{ij}_{\alpha\beta}(l)$ where i,j vary over labels of bins and α,β vary over cosmological probes such as galaxy number counts or shear measurements. (Indices may be suppressed when they are not relevant.) $C^{ij}_{\alpha\beta}(l)$ is symmetric in i,j and α,β.

2.1 The C(l) formalism

Following the approach of Peebles (1973) – and later Blake et al. (2004), Thomas et al. (2012) – for a particular probe of our cosmology observed projected on the sky in the direction of a unit vector $\underline{\mathrm{n}}$, $X(\underline{\mathrm{n}}) = \bar{X} + \Delta X(\underline{\mathrm{n}})$, we may decompose the variation in this parameter ΔX into spherical harmonics as

$$\Delta X(\underline{\mathrm{n}}) = \sum_{l>0} \sum_{m=-l}^{l} a_{lm} Y_{lm}(\underline{\mathrm{n}}) \tag{1}$$

We may calculate the coefficients a_{lm} by using the orthogonality of spherical harmonics ($\int Y_{lm} Y^*_{l'm'} d\Omega = \delta_{ll'}\delta_{mm'}$) :

$$a_{lm} = \int \Delta X(\underline{\mathrm{n}}) Y^*_{lm}(\underline{\mathrm{n}}) d\Omega \tag{2}$$

The $C(l)$s are defined from these coefficients by the relation:

$$C(l) = \langle a_{lm} a^*_{lm} \rangle \tag{3}$$

In our case we are interested in the galaxy distrtibution as a tracer of matter; this is calculated from the data by analysing number counts across the sky. It is important to note that this does not require knowledge of $n(z)$: we do not use redshift information in calculating the angular power spectrum from the data.

For the theoretical modelling however, we do require knowledge of $n(z)$, as we must calculate the full power spectrum $P(k,z)$ which is then projected onto the sky. This projection, as we shall later see, is strongly dependent on z. To calculate the $C(l)$s we use the following equation (Thomas et al. 2012):

$$C^{ij}_{\alpha\beta}(l) = \frac{2}{\pi} \int W^i_\alpha(l,k) W^j_\beta(l,k) k^2 P(k) dk \tag{4}$$

although we will perform most calculations without this approximation, it is useful to understand the impact of the redshift distributions. The redshift distributions $n^i(z)$ enter the $C(l)$s through the window functions.

2.2 Window functions

Window functions allow us to project the distribution of galaxies onto the sphere and decompose into spherical harmonics. Here we will discuss only the window function for galaxy clustering, since we have not used other probes in this particular work.

2.2.1 The galaxy clustering window function

Huterer et al. (2001) derive a calculation for galaxy clustering information of the following form.

$$W^i_{\rm g}(l,k,z) = \int b_g(k,z) n^i(z) j_l(k\chi) D(z) \tag{5}$$

Here $n^i(z)$ is the redshift distribution in bin i, $b_g(k,z)$ is the galaxy bias, $D(z)$ is the growth function, and $j_l(k\chi)$ is the order l spherical bessel function. Note that the comoving distance to an object is a function of redshift $\chi(z)$.

2.2.2 Including Redshift Space Distortions

Redshift Space Distortions (RSD) are alterations to the redshift of a galaxy due to its peculiar velocity. This leads to a distortion of the galaxy distribution if we attempt to reconstruct the three dimensional information, with galaxies with peculiar velocity toward us appearing closer (at lower redshift) and galaxies with peculiar velocity away from us along the line of sight appearing further (at higher redshift). Since these peculiar motions are due to interactions with local gravitational potentials they contain cosmological information. RSD on linear scales can be included by an additional term in the window function, following Kirk et al. (2013).

$$W^i_{\rm RSD}(l,k,z) = \beta \int \phi(\chi)[\frac{2l^2+2l-1}{(2l+3)(2l-1)} j_l(k\chi) - \frac{l(l-1)}{(2l-1)(2l+1)} j_{l-2}(k\chi) - \frac{(l+1)(l+2)}{(2l+1)(2l+3)} j_{l+2}(k\chi)] d\chi \tag{6}$$

The complete window function to be used in our $C(l)$ calculation is then given by the sum of these two terms.

$$W^i_{\rm LSS} = W^i_g(l,k) + W^i_{\rm RSD}(l,k) \tag{7}$$

In this paper we will not consider the effects of galaxy bias, although some papers have noted the potential importance of evolving galaxy bias in determining redshifts from correlation data (Schmidt (2006), Schulz (2015)). In the absence of a compelling bias model however, bias may be best handled as nuisance parameter (or parameters) which is also marginalised over. This is demonstrated in Clerkin et al. (2014), but for simplicity we choose a constant bias $b_g = 1$.

We can now see how the redshift distribution enters into the $C(l)$ formalism. If we have an accurate redshift distribution, for instance from a spectroscopic survey, then this is all we need to begin calculating our theoretical correlations. Unfortunately, photometric estimates are far from perfect, and photometric redshift errors if left unignored may produce unforeseen effects in our computed $C(l)$s. We will now seek to understand what some of these effects may be.

2.3 The significance of n(z)

Intuitively one might expect that number counts on the same patch of sky will be highly correlated when close in redshift and less correlated when widely separated. If we have a spectroscopic sample and a photometric sample that overlap in redshift, then they will contain objects in the same larger clustering structures, which we will be able to see as boosts in their correlations. We do not expect to see clustering over very large distances, so we expect that samples widely separated in redshift will show very weak cross-correlations.

We can put this intuitive understanding on a more mathematical foundation. From the definition of the window functions and the C(l) calculation, we can see how we expect $n(z)$ to affect our calculated $C^{ij}(l)$. If, for the sake of simplicity, we assume a k-independent bias $b_g(z)$, then the k dependence of $W(l,k)$ comes entirely from the spherical Bessel function $j_l(k\chi)$. The window functions oscillate as a function of k, made of contributions with different frequencies set by the spherical bessel functions in the integral. Hence the redshift range of the integral sets the range of frequencies present in the window function. If the distribution for a particular bin $n^i(z)$ is close to zero outside a particular range (for instance, if we model $n(z)$ as a top hat or Gaussian function) then the integral over z has a fairly small range which contributes significantly. If two bins $n^i(z)$ and $n^j(z)$ are separted in z by significantly more than their variance, then our two window functions $W^i(l,k)$ and $W^j(l,k)$ will have only very small contributions with the same frequency. The product of two oscillating functions with different frequencies will tend to average to zero when integrated over, so we would expect that the integral over these two window functions to be small. If, however, the redshift ranges overlap in regions of significant number density, then there will be significant contributions to both window functions with the same frequency and forms. These, when integrated over, will not average to zero and give a large contribution to $C^{ij}(l)$. Hence we expect the $C^{ij}(l)$s to be dependent on the amount of overlap between distribution functions in different bins, with significant overlaps in areas with high number density giving the strongest signals.

In addition to the overlap between bins, the spherical bessel function in equation 5 also tells us more about the redshift dependence of $C^{ij}(l)$: we expect stronger signals from distributions at lower redshift where the amplitude of $j_l(k\chi(z))$ is higher. So whilst the overlap between bins will determine the relative power in cross-correlations compared to auto-correlations, moving all the bins together up or down in redshift can shift the amplitudes of all the signals together. The redshift distribution is of course not the only thing which will affect our signal, and cosmological effects enter into our equations through the growth function ($D(z)$ in equation 5) and the power spectrum ($P(k)$ in equation 4). This is the source of a very important degeneracy between our redshift distributions and comsological parameters, particularly those such as A_s or σ_8 which strongly control the amplitude of $P(k)$.

In order to fix the redshift distribution, we need bins which overlap our photometric redshifts but are strongly anchored so that any changes in photometric bins, even moving coherently, will be captured by the $C(l)$s. For this we require spectroscopic data, which is well known enough to

have rigidly fixed $n(z)$, which overlaps our photometric bins. Note that our only criterion here is that our spectroscopic and photometric data overlap, and not - unlike with template and machine learning techniques - that the spectroscopic sample be an unbiased representation of the photometric sample.

3 MODELLING THE REDSHIFT DISTRIBUTIONS

3.1 Photmetric redshifts

Due to the uncertainty in photometric redshifts, we cannot obtain an accurate redshift for each object that we have in our sample. We instead model a bin as a broader function which captures the distribution of redshifts which would be binned together.

In this work we will model photometric redshift distributions within a bin as a Gaussian distribution defined by their mean and variance.

$$n(z,\mu,\sigma) = G(z,\mu,\sigma) = \left(\frac{1}{2\pi\sigma^2}\right)^{\frac{1}{2}} \exp\left[-\frac{(z-\mu)^2}{2\sigma^2}\right] \quad (8)$$

In order to model the uncertainties in $n(z)$, we need to be able to control the shape of the function in a quantitive way, ideally with as few parameters as possible. Each time we add a parameter, we are adding N_{bins} new dimensions to our parameter space to be explored by the MCMC and hence our computation becomes exponentially more expensive. The most important parameters are the mean and width of the distribution; other adjustments to the shape can be abandoned without too much impact but nevertheless the method is general and in wider, strongly non-Gaussian redshift bins higher moments may be taken into account if necessary. When using the mean and variance of a gaussian distribution we may adjust μ and σ directly using the analytic formula for a Gaussian. To vary a general distribution, or to change the shape in other ways, you may refer to the appendix.

Although this is the template used for all the bins in this study, we may also apply non-Gaussian distortions to these distributions to model more complex effects. It is also important to note that this method is by no means limited to Gaussian functions, and these may be easily replaced by an arbitrary function (with some parametrisation) $F(z,\mathbf{p})$, where $\mathbf{p}$ is the parameter vector to be marginalised over. In the simple example above $\mathbf{p} = (\mu,\sigma)$, although we may extend this to include skew and kurtosis, and have $\mathbf{p} = (\mu,\sigma,s,k)$ or some other vector of parameters. In a typical survey such as DES, photometric redshift bins have a standard deviation of approximately 0.1 to 0.2 (Lahav et al. 2009); for our purposes we will take $\sigma = 0.1$.

3.2 Spectroscopic redshifts

In order to include spectroscopic redshifts into the same formalism, we model spectroscopic redshifts in bins with much narrower distributions. This will take the form of a much narrower Gaussian. (A narrow top-hat function may also be used, but smooth continuous functions are often computationally more stable.) We assume that spectroscopic information is known well enough that we do not vary these bins in the same way as the photometry, and so there is no parameter vector $\mathbf{p}$ to marginalise over. Spectroscopic bins will be modelled with a width of $\sigma = 0.025$, which requires spectroscopic redshifts to be estimated to within a few percent. Their thickness may be determined by the nature and quality of the spectroscopic sample, or as a compromise with computational efficiency. A small number of wider bins is less computationally expensive than many narrow bins; the width of spectroscopic bins makes little to no impact on the length of computation (integrations are performed between fixed redshifts), however each additional bin adds two new parameters, which means that we have more integrations to perform (scaling as $N_{\rm bins}^2$), larger covariance matrices, and a much larger parameter space which scales exponentially in volume with the number of parameters. Narrow bins allow us to look at very localised correlations at the cost of this additional computation.

4 THE LIKELIHOOD FUNCTION AND MCMC METHODS

4.1 The likelihood function for C(l)s

The likelihood is calculated from $C(l)$ for each model (i.e. each parameter set) compared to the $C(l)$ calculated from the fiducial model. Let the fiducial model be known as model A, and the model we wish to investigate model B; we calculate a log-likelihood of seeing some fluctuations a_{lm} in the model B compared to the model A, and take then take an expectation value assuming the fiducial model A in the absence of any data, as described in Bucher et al. (2002). Since a_{lm} are stochastically generated, any given cosmology may generate a wide variety of a_{lm}, and each set of a_{lm} can give therefore give a different likelihood when compared against a model. Hence, with no reason to generate one particular set over another, one calculates the expectation value of these possible likelihoods, on the assumption that our a_{lm} were generated by the cosmology represented by A. This quantity is dependent only on the $C(l)$s calculated in each model, and the properties of the survey such as sky coverage and noise which remain constant throughout.

For measured a_{lm}, using the fact that the expectation value is zero, we have for a given cosmology X the relation

$$\mathrm{Var}(a_{lm}) = \langle |a_{lm}|^2 \rangle_X = C_X(l) + N(l) \quad (9)$$

where noise is assumed to be isotropic and uncorrelated (shot noise) and taken into account by the noise function $N(l)$. Assuming Gaussian distributions, we then have

$$P(a_{lm}|X) = \left(\frac{1}{2\pi(C_X(l)+N(l))}\right)^{\frac{1}{2}} \exp\left[-\frac{|a_{lm}|^2}{2(C_X(l)+N(l))}\right] \quad (10)$$

We wish to calculate the (expected) likelihood function

$$\langle L \rangle = \left\langle \log\left[\frac{P(a_{lm}|B)}{P(a_{lm}|A)}\right]\right\rangle_A \quad (11)$$

Given equation 10 we can write

$$\frac{P(a_{lm}|B)}{P(a_{lm}|A)} = \left[\frac{C_A(l)+N(l)}{C_B(l)+N(l)}\right]^{\frac{1}{2}} \exp\left[\frac{|a_{lm}|^2}{2(C_A(l)+N(l))} - \frac{|a_{lm}|^2}{2(C_B(l)+N(l))}\right] \quad (12)$$

Taking logs we obtain

$$\log\left(\frac{P_A}{P_B}\right) = \frac{1}{2}\log\left(\frac{C_A(l)+N(l)}{C_B(l)+N(l)}\right) + \frac{|a_{lm}|^2}{2(C_A(l)+N(l))} - \frac{|a_{lm}|^2}{2(C_B(l)+N(l))} \quad (13)$$

We then take the expectation value assuming A using equation 9

$$L = \frac{1}{2}\left[1 - \frac{C_A(l)+N(l)}{C_B(l)+N(l)} + \log\left(\frac{C_A(l)+N(l)}{C_B(l)+N(l)}\right)\right] \quad (14)$$

We must then take into account all of the a_{lm}, bearing in mind this expression is not dependent on m, so we have $2l+1$ identical terms for each l, and taking into account the fraction of the sky f_s observed.

$$L = \frac{f_s}{2}\sum_{l=2}^{l_{max}}(2l+1)\left[1 - \frac{C_A(l)+N(l)}{C_B(l)+N(l)} + \log\left(\frac{C_A(l)+N(l)}{C_B(l)+N(l)}\right)\right] \quad (15)$$

This is the likelihood we will use for CMB temperature information ($C^{TT}(l)$) or when we calculate the autocorrelation of a bin. When we have multiple bins or cosmological probes where cross correlations must be taken into account, then we will have more than one $C(l)$ function for each cosmology. We then a multivariate Gaussian distribution instead of simply the product of independent Gaussians. The covariance matrices are

$$[M_{X,l}]_{i,j} = C^{ij}_{(X)}(l) + \delta^{ij}N^i(l) \quad (16)$$

where X may be either A or B (with the relavent $C(l)$s calculated in the left hand side) and where $N(l)$ is the noise associated with the experiment. Noise is only added on the diagonal as shot noise between bins should not be correlated and hence not contribute to the covariance. This gives a probability distribution

$$P(\mathbf{a}_{lm}|X) = \frac{1}{((2\pi)^k|\mathbf{M}_{X,l}|)^{\frac{1}{2}}}\exp\left[-\frac{1}{2}\mathbf{a}^T_{lm}\mathbf{M}^{-1}_{X,l}\mathbf{a}_{lm}\right] \quad (17)$$

for a $k\times k$ covariance matrix (i.e. cross correlating k bins). Repeating the above analysis, and using the following relations (where l, m subscripts have been suppressed for clarity, and we use summation convention over i, j)

$$\left\langle \mathbf{a}^T\mathbf{M}^{-1}_X\mathbf{a}\right\rangle_A = \left\langle a_i M^{-1}_{X,ij}a_j\right\rangle_A = M_{A,ij}M^{-1}_{X,ij} = M_{A,ji}M^{-1}_{X,ij} = \left[\mathbf{M}_A\mathbf{M}^{-1}_X\right]_{jj} = \mathrm{Tr}\left[\mathbf{M}_A\mathbf{M}^{-1}_X\right] \quad (18)$$

we arrive at the analagous log-likelihood to equation 15 for multiple $C(l)$s

$$L = \frac{f_s}{2}\sum_l(2l+1)\left[\mathrm{Tr}\left(I - M_{A,l}M^{-1}_{B,l}\right) + \ln\left(\det\left(M_{A,l}M^{-1}_{B,l}\right)\right)\right] \quad (19)$$

It is easy to see that this is zero for A=B.

4.2 Noise parameters and survey assumptions

4.2.1 Galaxy number counts

We limit our model to shot noise, which is described by the noise function

$$N^i(l) = (\sigma^i(l))^2 = \frac{1}{\bar{n}^i} = \frac{4\pi f_s}{f^i_g N_g} \quad (20)$$

Table 1. Parameters for a Planck like CMB survey.

Band Frequency	70	100	143	217
Beam Width θ_b / arcsec	14.0	9.5	7.1	5.0
Noise Effective Temperature / $\mu K\sqrt{s}$	212	56	56	84
Detector Number $n_{\rm det}$	12	8	12	12

Where f_s is the fraction of the sky observed by the survey, f^i_g is the fraction of the total number of galaxies observed which lie within that redshift bin $n^i(z)$, and N_g is the total number of galaxies observed over the entire survey. For a DES like survery we assume that $N_g = 3\times10^8$, $f_s = 0.12$ (from $A = 5000\mathrm{deg}^2$), and a redshift range $0 < z \leqslant 2$ (Lahav et al. 2009).

4.2.2 CMB TT information

Here we have a slightly more complex function which must take into account more survey information. We base our parameters on a Planck-like survey, based on the parameters described in Abdalla & Rawlings (2008). Our noise function is

$$N^2_l = \sum_{chan}\frac{1}{(\sigma_c\theta_b)^2}\exp\left(-\frac{l(l+1)\theta_b}{8\ln 2}\right) \quad (21)$$

$$\sigma_c = \frac{T_{\rm NE}\theta_{\rm sky}}{\sqrt{n_{\rm det}t}\theta_b} \quad (22)$$

where θ_b is the beam width, $T_{\rm NE}$ is the noise effective temperature, $n_{\rm det}$ is the number of detectors, and t is the integration time assumed to be one year. We assume information is collected in four bands with parameters detailed in Table 1. For CMB information we assume $f_{\rm sky} = 0.65$.

4.3 Computational details: UCLCl and PLINY codes

$C(l)$ calculations are performed using the UCLCl code developed at UCL, and the CLASS Boltzmann code (Blas et al. 2011) for the generation of the primordial power spectrum and transfer function. Within UCLCl most functions, including $n(z)$, are represented using splines. The spline representation is advantageous for this work because it allows us to easily manipulate and deform $n(z)$ in non-linear ways without having to define an analytic function with some parametrisation. (This means we could take an arbitrary form from, for instance, data and still manipulate it in the way described in this paper.) We can vary the mean and variance for an arbitary distribution in a precise way. For the higher moments such as skew and kurtosis, we must vary these more heuristically for a general distribution, and these transformations may affect other moments. These may all be varied by applying transformations to the z variable of the $n(z)$ spline, as discussed in the appendix.

The MCMC analysis is performed using PLINY, a nested sampler designed for parallel computation. It calculates a chain of points in the parameter space, with likelihoods and prior weights, and also outputs an evidence calculation. In order to calculate the posterior weight for each

point in the chain we need to use Bayes' Theorem:

$$\text{Posterior} = \frac{\text{Likelihood} \times \text{Prior}}{\text{Evidence}} \tag{23}$$

For all parameters in this analysis we assume flat priors with hard edges well away from the peak of the distribution. The evidence is not strictly necessary in this analysis as it is just a constant factor. The evidence is only required if we wish to perform a model comparison for models with different parameterisations.

4.4 The fiducial model

In this work we use a fiducial ΛCDM model. For the sake of computational efficiency, we take work only with flat cosmologies ($\Omega_k = 0$). We also restrict ourselves to varying seven cosmological parameters - $\{A_s, \Omega_\Lambda, \Omega_b, h, n_s, \tau_r, w_0\}$. Our fiducial cosmology will be

$$\begin{aligned} A_s = 25 \times 10^{-10}, \\ \Omega_\Lambda = 0.7, \\ \Omega_b = 0.06, \\ h = 0.7, \\ n_s = 0.95, \\ \tau_r = 0.09, \\ w_0 = -0.9 \end{aligned} \tag{24}$$

To speed computation we limit ourselves to flat cosmologies, and hence we will use $\Omega_{\rm cdm} = 1 - \Omega_\Lambda - \Omega_b$, which gives a fiducial $\Omega_{\rm cdm} = 0.24$. We use three photometric bins with mean

$$(\mu_1, \mu_2, \mu_3) = (0.8, 1.0, 1.2)$$

and standard deviation

$$\sigma_1 = \sigma_2 = \sigma_3 = 0.1$$

For analyses with spectroscopy, we use three spectroscopic bins with mean

$$(\mu_1, \mu_2, \mu_3) = (0.7, 1.0, 1.3)$$

and standard deviation

$$\sigma_1 = \sigma_2 = \sigma_3 = 0.025$$

Spectroscopic bins are assumed to be well known enough not to need variation in the MCMC analysis, so this leaves us with an 13-dimensional parameter space (7 cosmological and 6 photometric binning parameters).

For our noise function, in this particular analysis, we have chosen DES like parameters (described in section 4.2) with $f_s = 0.12$, $f_g^i = 0.2$ for all photometric bins, and $N_g = 3 \times 10^8$ (giving 6×10^7 galaxies in each photometric bin). For spectroscopic bins, we assume that we have 5×10^4 galaxies in each bin, roughly in keeping with the density in surveys such as BOSS or eBOSS (SDSS Collaboration 2008, www.sdss.org).

For the sake of reasonably rapid calculations we use a limited number of redshift bins in this demonstration, although most full surveys will use 5-10 photometric redshift bins. This will diminish our power to constrain the cosmological parameters somewhat due to a lack of information and coverage over much of the redshift range, but will be enough to demonstrate the power of the technique applied to the calibration of photometric redshifts.

5 RESULTS

In this section we will present the results of our MCMC analysis. We will first demonstrate the bias in cosmological parameters that is caused by having poorly estimated photo-z bins. We will then show what can be achieved using autocorrelations of photometric bins, where bins are allowed to vary freely; this will demonstrate where the degeneracies between photometric redshifts and cosmological parameters lie. Finally we will show results using cross correlations between both photometric and spectroscopic bins, which gives dramatically improved precision on the photometric bins, and we demonstrate the effect of this on the marginalised distributions for the cosmological parameters.

5.1 Cosmological parameter bias from n(z)

From equations 4 & 5, our theoretical prediction of $C^{ij}(l)$ is dependent on the redshift distributions $n^i(z)$ and $n^j(z)$. If we estimate properties of our redshift bins (in this case μ, σ) by fitting objects with known spectroscopic redshifts, then we will derive redshifts with some scatter around their 'true' value. These redshift errors have a knock-on effect on our inference of cosmological parameters. For example, if our estimated redshifts are too low, then theoretical power that we calculate will be too high; in order to match the observations, A_s may be lowered to match the power, and other parameters adjusted to get the best fit to shape. In this section we will demonstrate such biases, and later we shall see how marginalising over redshift distributions can avoid them. Methods in estimating photometric redshifts often have an error in z of $O(0.1)$, which is large compared to what we would require to obtain precise results from a photometric survey. If take the parametrisation obtained from this fitting on face value then we will reconstruct a slightly distorted $n(z)$. This means that when we fit our cosmological model, our cosmological parameters will be inevitably be changed in order to counter the effect of the distortions in $n(z)$. We may analyse this case in our simple model by using the fiducial $n(z)$ for the "observed" $C(l)$s as described in section 4.4, but calculating our model $C(l)$s using bins fixed to have different parameters.

The cosmology used is the same as stated in section 4.4 but we shall only use two redshift bins at $\mu_1 = 0.8$ and $\mu_2 = 0.9$ to generate the fiducial $C(l)$s. When we attempt to recover the cosmological parameters with an MCMC analysis, we use a fixed redshift distribution, biased with $\mu_1 = 0.75$ and $\mu_2 = 0.85$. The results are shown in Figure 1.

The bias is strongest in cases such as this where there is a systematic error causing the mean or standard deviation of bins to be consistently over or under estimated. In order to avoid this, we must reduce our reliance on fixed redshift distributions with large errors. In lieu of a method for sufficiently accurate redshifts from photometry, we must rely on marginalising in a bayesian framework, the results of which are described in the following sections.

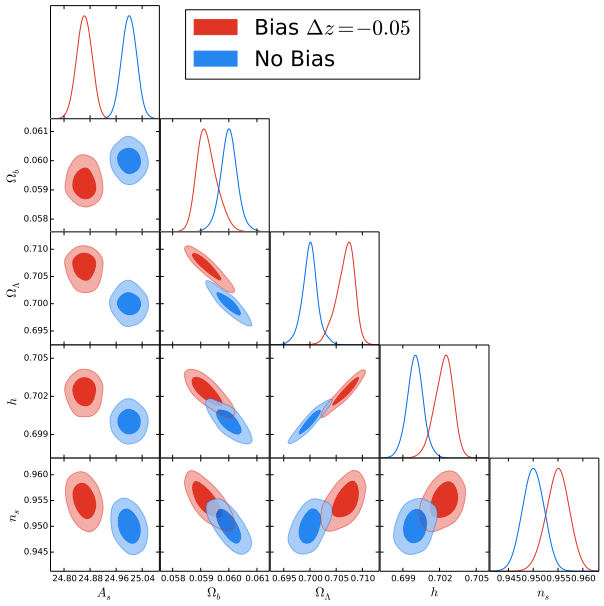

Figure 1. Probability contours obtained from the true redshift distribution (shown in blue) and from a biased redshift distribution (shown in red). A five parameter cosmology is derived from two photometric bins, where the fiducial cosmology and the blue contours use $\mu_1 = 0.8$ and $\mu_2 = 0.9$, whereas the red contours are derived on the incorrect assumption that $\mu_1 = 0.75$ and $\mu_2 = 0.85$ i.e. photometric redshifts are systematically underestimated. The blue contours are, by construction, centred on the fiducial parameters, whereas the red contours end up far from the fiducial parameters in order to compensate for effects in the $C(l)$ signal introduced by photometric systematics.

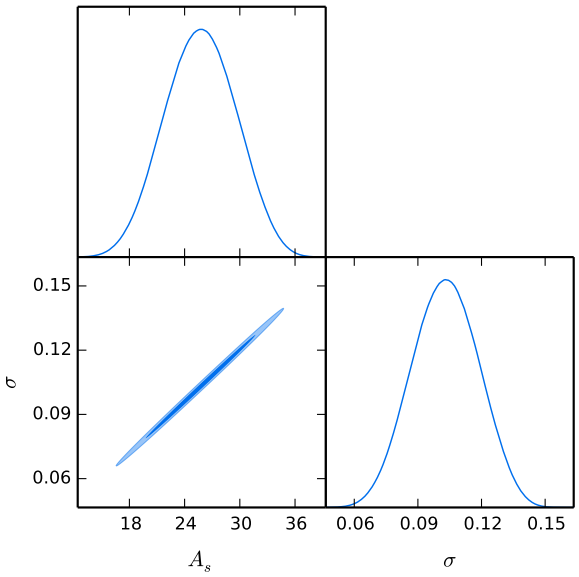

Figure 2. Probability contour obtained when varying only A_s and σ (the width of one photometric redshift bin), with $\mu = 1.0$ and all other parameters at their fiducial values, demonstrating the high degeneracy between photometric bin width and the A_s parameter.

5.2 Autocorrelations with photometric redshift bins

The simplest analysis that we can do is to use only our photometric redshift bins, and to only take into account autocorrelations. We will see that this means ignoring a great deal of information, and our bounds on cosmological and binning parameters are wide. Although in this case we are not taking into account the full information available to us, it is worth looking into since it is much less computationally expensive, and previous studies have been focussed on autocorrelations. We promote the mean and standard deviation of our redshift bins to fully independent parameters for our MCMC analysis, allowing them to vary freely so that they can be marginalised over. Here we use the three photometric bins described in section 4.4. This information is combined with CMB TT information in order to help constrain A_s, which is a problematic parameter in this analysis since it is extremely strongly degenerate with the standard deviation of a bin (see Figure 2).

Despite the fact that we are not using any cross-correlations between bins, we can see the parameters for different bins are degenerate. This is because of the effect that μ and σ have on the $C(l)$. For example, when all the bins are moved in the same direction, the effect is largely to raise or lower the power in each autocorrelation; this can be compensated for by adjusting A_s and other cosmological parameters appropriately. If however some bins are moved up in redshift, and some down, then the cosmological parameters struggle to compensate for the competing effects. The same is true for σ. This means that these parameters are constrained to move together to some extent.

We note that the degeneracy between σ and A_s is being prevented from exercising its full effect because A_s has been constrained significantly by the CMB information. Nevertheless A_s is strongly degenerate with τ_r (the optical depth at reionisation) and we note that τ_r is not well constrained in this instance. Although its effects have been mitigated by its constraint, it is still clear that there is a degeneracy between σ of each bin and all of the cosmological parameters except n_s. Likewise μ is strongly degenerate with Ω_b, Ω_Λ, h, and w_0, and is not constrained up to the hard limits of the prior. This means that errors in estimates for binning parameters can propagate into cosmological parameters in a significant way.

We can understand the degeneracy between the cosmological parameters by considering their effects upon the $C(l)$s. The effect of σ is primarily to change the height of the $C(l)$s, which creates its degeneracy with A_s. Likewise, we know that moving μ to low redshift boosts power; since Ω_Λ and h suppress structure formation, these need to be lowered and Ω_b raised to get the C(l)s to match the fiducial model.

5.3 Cross correlating with spectroscopic redshifts

In this section we demonstrate the improvement attainable by cross-correlating with spectroscopic redshift data. Because of the overlap with spectroscopic data, we can show that the properties of the photometric redshift bins are now tightly constrained, and the degeneracies between bins are less pronounced. In most cases, the binning parameters cannot vary widely enough to have a noticeable impact on the

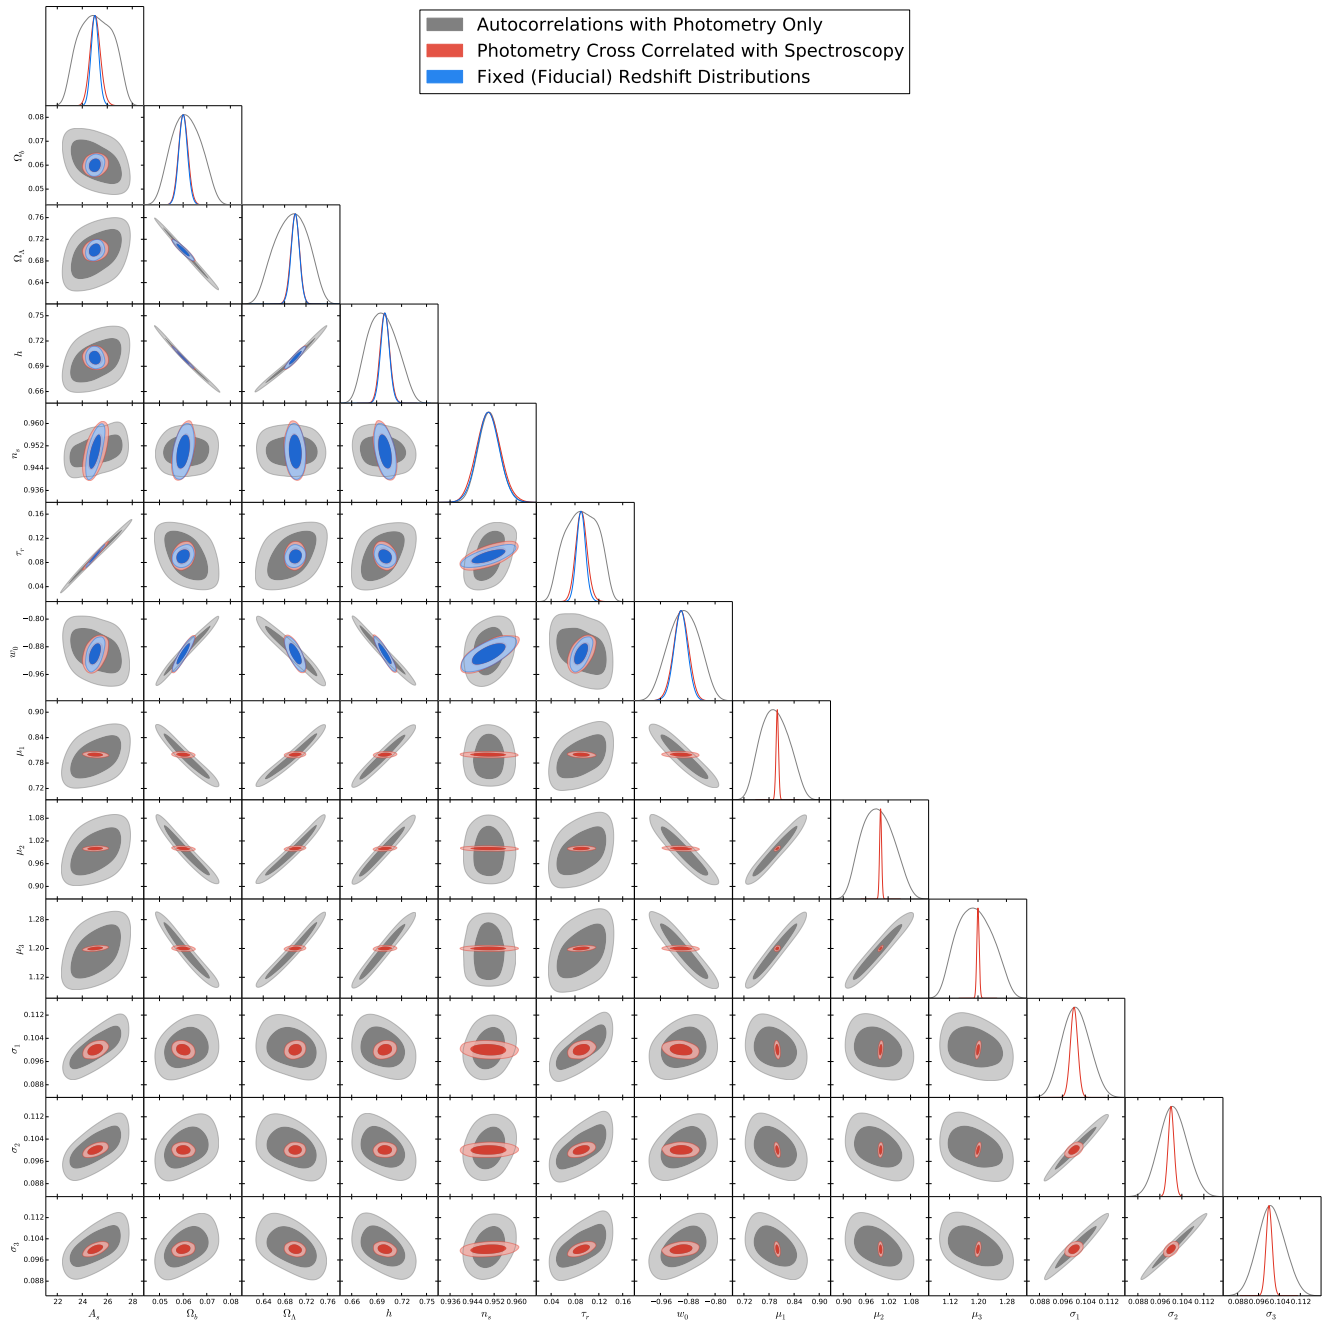

Figure 3. Probability contours obtained for cosmological and photometric binning parameters. Results from using autocorrelations of photometric bins only are shown in grey, and results from using photometric bins cross correlated with each other and spectroscopic bins are shown in red. For comparison, results using redshift bins fixed at the fiducial values are shown in blue. Results show clear improvements on all the binning parameters, as well as most cosmological parameters (with the exception of n_s), and contours using cross correlations between photometric and spectroscopic samples yield results very close to those with no redshift error.

cosmological parameters compared to the uncertainty already present.

We can see that there is increased precision in the cosmological parameters (except for n_s, which is almost entirely determined by CMB information here), with most bounds improving by a factor of two or more (Table 2). The slight widening in the posterior distribution for n_s which can be seen in Figure 3 is most likely due to the additional noise introduced to the galaxy clustering likelihood by looking at larger numbers of bins. Since galaxy number counts do little to constrain n_s at this level, and bins which are widely separated in redshift may produce correlation functions that are largely noise dominated (since they should be close to zero), this small additional of noise to the likelihood causes some spreading of this parameter. This could be tackled by ignoring widely separated bins if necessary (this would also speed the likelihood calculation by reducing the number of integrations). Binning parameters (μ_i, σ_i) have been partic-

Table 2. 68% confidence ranges for inferred cosmological parameters using only autocorrelations, and using cross correlations with spectroscopy.

	$A_s \times 10^{10}$	Ω_b	Ω_w	w_0	h	τ_r	n_s
Fiducial	25	0.06	0.7	-0.9	0.7	0.09	0.95
Photometric	25.1 ± 1.5	0.060 ± 0.007	0.693 ± 0.033	-0.89 ± 0.05	0.697 ± 0.019	0.091 ± 0.029	0.950 ± 0.004
Photo × Spec	25.0 ± 0.4	0.060 ± 0.002	0.700 ± 0.008	-0.90 ± 0.02	0.700 ± 0.006	0.091 ± 0.009	0.950 ± 0.004
Fixed Redshift	25.0 ± 0.3	0.060 ± 0.002	0.700 ± 0.008	-0.90 ± 0.02	0.700 ± 0.005	0.090 ± 0.008	0.950 ± 0.004

Table 3. 68% confidence ranges for inferred photometric redshift bin parameters, using only autocorrelations and using cross correlations with spectroscopy.

	μ_1	μ_2	μ_3	σ_1	σ_2	σ_3
Fiducial	0.8	1.0	1.2	0.1	0.1	0.1
Photo+Auto	0.794 ± 0.036	0.992 ± 0.045	1.191 ± 0.054	0.101 ± 0.005	0.101 ± 0.005	0.101 ± 0.005
Spec+Cross	0.800 ± 0.003	1.000 ± 0.003	1.200 ± 0.003	0.100 ± 0.001	0.100 ± 0.001	0.100 ± 0.001

ularly tightly constrained, allowing us to know their values to percent level or better (Table 3). Further constraint can be imposed upon them by having more spectroscopic bins to cover a greater fraction of the photometric redshift range, at the cost of computation time.

5.4 Constraining higher moments

As a benchmark, we also present constraints on the shapes of a single photometric bin from cross correlations with three spectroscopic bins. Here we have a fixed cosmology to simplify the calculations and provide benchmark results for the shape parameters. With cross correlations between more photometric and spectroscopic bins we expect these results to be improved.

We find that the odd moments are strongly correlated, but we also have less information to constrain the higher moments from the correlation statistics. This may not be a problem unless the higher moments significantly affect the cosmological parameter estimation, although the lack of constraining power of the $C(l)$s suggest that these are not strongly affected by the finer details of the shape of the distributions. The most significant potential problem is the degeneracy between s and μ. This could lead to a spreading of the distribution over mu which may affect the cosmological parameter estimation. In order to combat this one would need to have relatively dense spectroscopic samples throughout the range. In principle, the degeneracy between μ and s need not be a problem. Since they are strongly degenerate, it is largely the case of the one compensating for the effect of the other. In this case, the cosmological calculation may not be strongly affected even if the uncertainty in μ increases significantly. We should also bear in mind that the application of the skew transformation alters the mean of the distribution even thought the peak is kept in the same place. This needs to be compensated for by the μ parameter, and thus the true mean may remain roughly unchanged and be very strongly constrained. We do not expect higher moments to have a very strong impact on cosmological results.

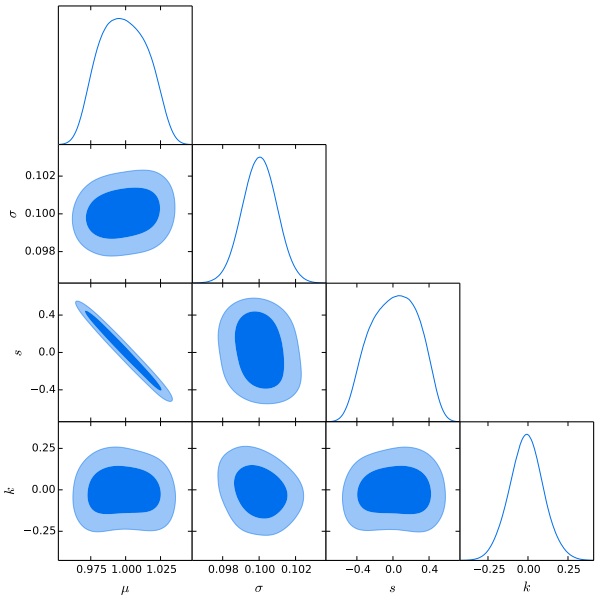

Figure 4. Constraints on the mean, width, and shapes of bins as characterised by (μ, σ, s, k) for a single photometric bin, cross correlated with three spectroscopic bins, with fixed cosmology.

6 CONCLUSIONS

From the results presented in section 5, we can see that different aspects of the analysis provide distinct benefits. The variation of $n(z)$, and its subsequent marginalisation, is essential for the removal of the bias from cosmological inferences. In order utilise the full power of the $C(l)$ formalism, we must include cross-correlations as well as the well studied auto-correlations; these not only provide us with with much more information (improving our constraining power), but also help us to pin down the relationships between different photometric redshift bins more accurately. Since the photometric redshift parameters display degeneracy with almost all of the cosmological parameters, it is crucial to have these distributions known as well as possible.

By using a full theoretical likelihood and MCMC approach this study is less idealised than most previous works; nevertheless the assumption of gaussian bins, and the small number of redshift bins, are simplifications that should be addressed in future work. As the analysis is extended with more bins, the computational complexity will increase significantly; the number of cross-correlations to calculate will increase as $O(n_{bins}^2)$ and the dimensionality of the parameter space as $O(n_{bins})$. It is possible to simplify in such cases by only considering cross-correlations between bins which are sufficiently close together (leading to a band diagonal $C^{ij}(l)$ matrix), since cross-correlations between widely separated bins will contain comparatively little information. Despite the computations intensity, as long as numerical errors remain tamed, we expect that increasing the number of bins and the density of the spectroscopic sample to improve the results. A survey such as Euclid should have a great deal of power to jointly constrain the redshift distribution and the cosmology with minimal disturbance to the confidence intervales for cosmological parameters.

Errors on redshift binning parameters are now $O(10^{-3})$, even with such sparse spectroscopic data as we have simulated. Errors on the means of photometric bins are at ± 0.003 and errors on the width of bins is at ± 0.001 (see Table 3). This is extremely promising for future experiments, providing the possibility to extract reliable and precise cosmological parameters. As we look towards future experiments such as Euclid and LSST, and even with data currently being released from DES, a major focus in cosmology will be the nature of dark energy. The ability to distinguish between a cosmological constant, scalar field theory, modified gravity, or more exotic forms still, will be dependent on having well known redshift distributions, as can be seen by the strong degeneracy between the mean of redshift bins and the parameters Ω_Λ and w_0. A bias from improperly calibrated photometric data could easily generate a spurious result. Using this method will help to ensure robust analyses for current and future experiments.

This method can be applied to any $C(l)$ signal using photometric redshift bins (such as weak lensing or galaxy clusters) to calibrate their photometric redshift distributions. This means that when applied to future optical surveys, it will be able to benefit much more powerful analyses than the one outlined in this paper, including a larger number of redshift bins, and a combination of signals from different cosmological probes. If the same photometry is used for both number counts and lensing, then both of these cross correlations will contribute to constraining the photometric parameters, as well as constraining the cosmology itself. It will be necessary for future observational work to extend this to non-gaussian distributions, including higher order moments or more generic spline models of $n(z)$, in order to model our observed photometric redshifts as best we can. When combined with lensing information, this technique can be applied to achieve improved results in Modified Gravity or Dark Energy studies, where biases can lead to spurious detections and high precision is needed.

ACKNOWLEDGEMENTS

MM is supported by an UCL studentship; FBA acknowledges the Royal Society for a Royal Society University Research Fellowship.

REFERENCES

Abdalla F., Rawlings S., 2008, MNRAS, 381, 1313
Abdalla F. B., Banerji M., Lahav O., Rashkov V., 2008, MNRAS, 417, 1891
Banerji M., Abdalla F., Lahav O., Lin H., 2008, MNRAS, 386, 1219
Benjamin J. et al, 2010, MNRAS, 408, 1168
Blake C., Ferreira P. G., Borrill J., 2004, MNRAS, 351, 923
Blake C. Collister A., Bridle S., Lahav O., 2007, MNRAS, 374, 1527
Blas D., Lesgourgues J., Tram T., 2011, JCAP, 07, 034
Bucher M., 2002, Phys. Rev. D 66, 023528
Clerkin L. et al., 2014, MNRAS 448, 1389
Huterer D., 2001, ApJ, 555, 547
Huterer D., et al., 2004, ApJ, 615, 595
Kirk D., Lahav O. et al., 2013, MNRAS, 438, 2218
Lahav O., Kiakotou A., Abdalla F. B., Blake C., 2009, MNRAS, 405, 168
Matthews D., Newman J., 2010, ApJ, 721, 456
McQuinn M., White M., 2013, MNRAS, 433, 2857
Menard B. et al., 2011, arXiv:1303.4722v2
Newman J., 2008, arXiv:0805.1409
Newman J. A. et al., 2012, Astroparticle Physics, 63, 81
Peebles P. J. E., 1973, ApJ, 185, 413
Rhodes J. et al., 2014, arXiv:1309.5388v3
Sadeh I., Abdalla F., Lahav O., 2015, arXiv:1507.00490v1
Schmidt S. et al., 2013, MNRAS, 431, 3307
Schulz A.E., 2006, ApJ, 724, 1305
SDSS Collaboration, 2008, Astron. J, 142, 72
Thomas S. A., Abdalla F., Lahav O., 2012, MNRAS, 412, 1169
Zheng H. & Zhang Y., 2012, Proc. SPIE 8451, Software and Cyberinfrastructure for Astronomy II, 845134

APPENDIX A: DETAILS OF N(Z) TRANSFORMATIONS

In this appendix we will describe in detail the transformations made to the $n(z)$ functions, and the motivations for the heuristic shape manipulation. It is not important that the higher moments are not exactly represented in the same way that the mean and variance of the gaussian are – these are after all only parameters controlling the shape which will be marginalised over. The important thing is that it can explore a variety of shapes with a small number of parameters. More precise handling of the distributions can be achieved at the cost of increasing the parameter space, which may rapidly make the computation unmanageable without abundant computing resources.

A1 Mean and Variance

Due to the spline representation used in our computation (see section 4.3) it is simplest to perform a transformation the z axis to a new variable $z' = f(z)$ for a general (non-Gaussian) distribution. By creating the function $n'(z') = n(z) = n(f^{-1}(z))$, we obtain a distorted distribution in our new variable, which we take to be the new redshift. (This distribution, as all distributions, is normalised before any further calculations are carried out.)

To vary the mean ($\mu \to \mu + \Delta\mu$), we apply the transformation:

$$z' = z + \Delta\mu \tag{A1}$$

This varies the mean of an arbitrary distribution without affecting any of the higher moments. We can also change the standard deviation ($\sigma \to \sigma + \Delta\sigma$) without affecting the mean or higher moments. (The fourth moment m_4 is changed, but not kurtosis $\kappa = \frac{m_4}{\sigma^4}$.)

$$z' = \left(1 + \frac{\Delta\sigma}{\sigma}\right)(z - \mu) + \mu \tag{A2}$$

One is free to change the mean and standard deviation of distributions in whichever order desired as these transformations are commutative. Higher order transforms may also be applied which break symmetries, so care must be taken there as the operations will not be commutative.

A2 Skewness

To adjust the apparent skewness, we need to stretch the distribution on one side of the mean, and squeeze the distribution on the other. For the sake of simplicity, we write a heuristic skew function controlled by a single parameter s. (This is to distinguish it from the skewness calculated from the third moment, γ.) The parameter range is defined at $-1 < s < 1$. We map from the original redshift coordinate z to a new coordinate z' representing the new redshift after the distortion has been taken into account. If the two are identical then we have $\frac{dz'}{dz} = 1$ everywhere. If we wish to stretch a region then $\frac{dz'}{dz} > 1$ and to squeeze it we have $\frac{dz'}{dz} < 1$. To achieve skewness, we need to smoothly vary from stretched regions on the one side of the mean, to squeezed regions on the other side, with $\frac{dz'}{dz}|_{z=\mu} = 1$.

We may choose a simple linear function:

$$\frac{dz'}{dz} = 1 + \frac{(z-\mu)s}{L} \tag{A3}$$

This fulfils the criteria discussed in the range $\mu - L < z < \mu + L$. After these points we fix $\frac{dz'}{dz}|_{z<\mu-L} = \frac{dz'}{dz}|_{z=\mu-L}$ and $\frac{dz'}{dz}|_{z>\mu+L} = \frac{dz'}{dz}|_{z=\mu+L}$. This is because otherwise we rapidly end up with very extreme stretching or squeezing of the distribution. Here we have an extra free parameter, L (the lengthscale of the skewness function). To avoid overburdening the routine with extra parameters we generically set this to $L = \frac{3\sigma}{4}$. The resulting function is found by integrating these expressions with the condition that $\mu' = \mu$. (This means that skewness does not interfere with the peak of the distribution, but will change the mean; it may also interfere with standard deviation. Standard deviation and mean can separately be readjusted to remove this degeneracy if desired.)

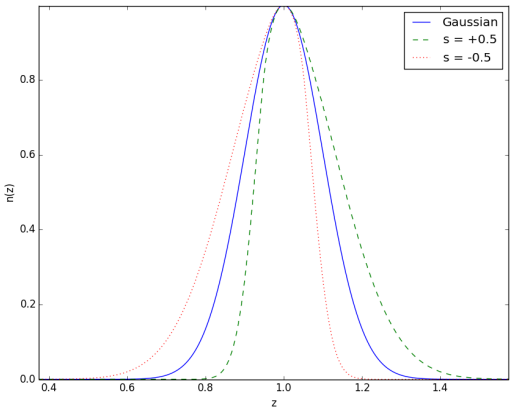

Figure A1. Results of applying positive and negative skew transformations to a gaussian photometric bin.

$$z' = \begin{cases} \frac{s}{L}\left(\frac{1}{2}(z+\mu)^2 - \mu\right) + z, & |z-\mu| \leqslant L \\ (\mu + L)' + (1+s)(z - (\mu + L)), & z - \mu > L \\ (\mu - L)' + (1-s)(z - (\mu - L)), & z - \mu < -L \end{cases} \tag{A4}$$

A3 Kurtosis

Kurtosis is handled in a similar way to skewness, by a heuristic function controlled by a single parameter k (different to the exact kurtosis, traditionally κ) which varies $-1 < k < 1$. This case is symmetric, and we wish to stretch the distribution close to the mean, and squeeze it further away from the mean (or vice versa).

In this case we again need to choose length scales. We choose to have the transition from stretched to squeezed regions (i.e. $\frac{dz'}{dz} = 1$) at $(z - \mu) = \sigma$, and then fix the transformations (as with the skewness above) at $(z - \mu) = 2\sigma$. This requires in the below expression $L = \sigma$.

For simplicitly we again choose linear relations.

$$\frac{dz'}{dz} = (1 + k) - \frac{|z - \mu|}{L}k \tag{A5}$$

We then integrate as before, choosing $\mu' = \mu$. Once again, the standard deviation can be separately adjusted for if desired.

$$z' = \begin{cases} z + k(z-\mu) - k\frac{(z-\mu)^2}{2\sigma}, & \mu \leqslant z \leqslant \mu + 2\sigma \\ z + k(z-\mu) + k\frac{(z-\mu)^2}{2\sigma}, & \mu - 2\sigma \leqslant z \leqslant \mu \\ \mu + \sigma + (1-k)(z - \mu - 2\sigma), & z \geqslant \mu + 2\sigma \\ \mu - \sigma + (1-k)(z - \mu + 2\sigma), & z \leqslant \mu - 2\sigma \end{cases} \tag{A6}$$

This paper has been typeset from a TEX/ LATEX file prepared by the author.

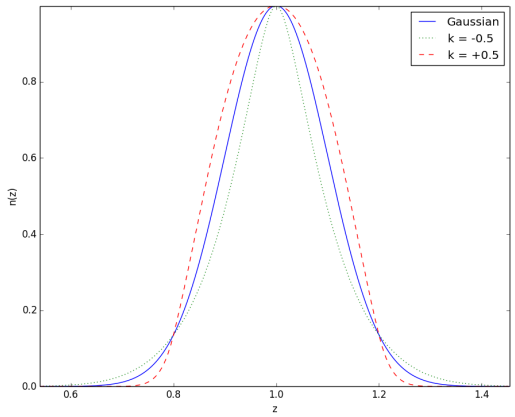

Figure A2. Results of applying positive and negative kurtosis transformations to a gaussian photometric bin.